

SYSTEM, METHOD, AND COMPUTER PROGRAM PRODUCT FOR SECURE REAL-TIME N-PARTY COMPUTATION

CROSS REFERENCE TO RELATED APPLICATION

[0001] This application claims priority to United States Provisional Patent Application No. 63/124,159, filed December 11, 2020, the disclosure of which is hereby incorporated by reference in its entirety.

BACKGROUND

1. Field

[0002] Disclosed embodiments or aspects relate generally to fair exchange computer protocols and, in one particular embodiment or aspect, to a system, method, and computer program product for secure real-time n-party computation.

2. Technical Considerations

[0003] Secure multiparty computation (MPC) allows a set of parties to perform a joint computation on their inputs that reveals only the final output and nothing else. These parties may be mutually distrusting and require a secure protocol to ensure a fair exchange. Secure computation, however, often cannot provide fairness in settings where many (e.g., a majority, a plurality, half, and/or the like) of the participants are corrupt. For instance, in a two-party setting, a malicious party can abort a secure computation protocol after receiving the output, leaving the other party no recourse to obtain the output. Additionally, use of a third party as an intermediary for such multiparty computations may waste resources in terms of the number of parties involved, the number of communications needed, and/or the like, and it may be possible for the third party to collude with one of the parties, thereby threatening security and fairness. There is a need to address these and other deficiencies of secure computation.

SUMMARY

[0004] Accordingly, and generally, provided is an improved system, method, and computer program product for secure real-time n-party computation.

[0005] According to some non-limiting embodiments or aspects, provided is a system for multi-party computation between a first party and a second party. The system includes a first computing device associated with the first party and a trusted

execution environment (TEE) associated with the first computing device. The first computing device is programmed or configured to communicate, to the TEE, a first computation input and a first portion of a one-time key. The first computing device is also programmed or configured to receive, from the TEE, an encrypted output of a computation based on the first computation input and a second computation input. The second computation input and a second portion of the one-time key are communicated to the TEE by a second computing device associated with the second party. The encrypted output is encrypted by the TEE with the one-time key. The first computing device is further programmed or configured to communicate the encrypted output to the second computing device and receive a digital signature from the second computing device indicating that the second computing device received the encrypted output. The first computing device is further programmed or configured to communicate the first portion of the one-time key to the second computing device. The first computing device is further programmed or configured to: (i) in response to receiving the second portion of the one-time key from the second computing device, determine the one-time key from the first portion of the one-time key and the second portion of the one-time key, and determine the computation by decrypting the encrypted output with the one-time key; or (ii) in response to not receiving the second portion of the one-time key from the second computing device, execute a fallback computation process using the TEE and a ledger accessible to both the first computing device and the second computing device to determine the computation.

[0006] In further non-limiting embodiments or aspects, the first computing device may be further programmed or configured to, in the fallback computation process: publish, to the ledger, the first portion of the one-time key and the digital signature received from the second computing device; receive a proof of publication from the ledger; input the proof of publication to the TEE; receive the one-time key from the TEE; and determine the computation by decrypting the encrypted output with the one-time key.

[0007] In further non-limiting embodiments or aspects, the first portion of the one-time key that is published to the ledger by the first computing device may be encrypted with a public key associated with the second computing device.

[0008] In further non-limiting embodiments or aspects, the one-time key may be generated by the TEE from the first portion of the one-time key, the second portion of the one-time key, and a unique nonce associated with an instance of the computation.

[0009] In further non-limiting embodiments or aspects, the first portion of the one-time key may be generated from a private key associated with the first computing device, and the second portion of the one-time key may be generated from a private key associated with the second computing device.

[0010] In further non-limiting embodiments or aspects, the first portion of the one-time key and the second portion of the one-time key may be communicated to the TEE in a first time period and may be configured to generate a plurality of one-time keys over successive instances of multi-party computation in subsequent time periods.

[0011] In further non-limiting embodiments or aspects, the ledger may include an append-only blockchain and may be configured to be accessible only by computing devices that contribute to the computation.

[0012] According to non-limiting embodiments or aspects, provided is a computer program product for multi-party computation between a first party and a second party. The computer program product is stored on a first computing device associated with the first party and includes at least one non-transitory computer-readable medium including program instructions. The program instructions, when executed by at least one processor of the first computing device, cause the at least one processor to communicate, to a trusted execution environment (TEE) associated with the first computing device, a first computation input and a first portion of a one-time key. The program instructions further cause the at least one processor to receive, from the TEE, an encrypted output of a computation based on the first computation input and a second computation input. The second computation input and a second portion of the one-time key are communicated to the TEE by a second computing device associated with the second party. The encrypted output is encrypted by the TEE with the one-time key. The program instructions further cause the at least one processor to communicate the encrypted output to the second computing device and receive a digital signature from the second computing device indicating that the second computing device received the encrypted output. The program instructions further cause the at least one processor to communicate the first portion of the one-time key to the second computing device. The program instructions further cause the at least one processor to: (i) in response to receiving the second portion of the one-time key from the second computing device, determine the one-time key from the first portion of the one-time key and the second portion of the one-time key, and determine the computation by decrypting the encrypted output with the one-time key; or (ii) in

response to not receiving the second portion of the one-time key from the second computing device, execute a fallback computation process using the TEE and a ledger accessible to both the first computing device and the second computing device to determine the computation.

[0013] In further non-limiting embodiments or aspects, the program instructions may further cause the at least one processor to, in the fallback computation process: publish, to the ledger, the first portion of the one-time key and the digital signature received from the second computing device; receive a proof of publication from the ledger; input the proof of publication to the TEE; receive the one-time key from the TEE; and determine the computation by decrypting the encrypted output with the one-time key.

[0014] In further non-limiting embodiments or aspects, the first portion of the one-time key that is published to the ledger by the at least one processor may be encrypted with a public key associated with the second computing device.

[0015] In further non-limiting embodiments or aspects, the one-time key may be generated by the TEE from the first portion of the one-time key, the second portion of the one-time key, and a unique nonce associated with an instance of the computation.

[0016] In further non-limiting embodiments or aspects, the first portion of the one-time key may be generated from a private key associated with the first computing device, and the second portion of the one-time key may be generated from a private key associated with the second computing device.

[0017] In further non-limiting embodiments or aspects, the first portion of the one-time key and the second portion of the one-time key may be communicated to the TEE in a first time period and may be configured to generate a plurality of one-time keys over successive instances of multi-party computation in subsequent time periods.

[0018] In further non-limiting embodiments or aspects, the ledger may include an append-only blockchain and may be configured to be accessible only by computing devices that contribute to the computation.

[0019] According to non-limiting embodiments or aspects, provided is a computer-implemented method for multi-party computation between a first party and a second party. The method includes communicating, with at least one processor of a first computing device associated with the first party, to a trusted execution environment (TEE) associated with the first computing device, a first computation input and a first portion of a one-time key. The method also includes receiving, with at least one

processor of the first computing device, from the TEE, an encrypted output of a computation based on the first computation input and a second computation input. The second computation input and a second portion of the one-time key are communicated to the TEE by a second computing device associated with the second party. The encrypted output is encrypted by the TEE with the one-time key. The method further includes communicating, with at least one processor of the first computing device, the encrypted output to the second computing device. The method further includes receiving, with at least one processor of the first computing device, a digital signature from the second computing device indicating that the second computing device received the encrypted output. The method further includes communicating, with at least one processor of the first computing device, the first portion of the one-time key to the second computing device. The method further includes, in response to not receiving the second portion of the one-time key from the second computing device, executing, with at least one processor of the first computing device, a fallback computation process using the TEE and a ledger accessible to both the first computing device and the second computing device to determine the computation.

[0020] In some non-limiting embodiments or aspects, the fallback computation process may include: publishing, with at least one processor of the first computing device, to the ledger, the first portion of the one-time key and the digital signature received from the second computing device; receiving, with at least one processor of the first computing device, a proof of publication from the ledger; inputting, with at least one processor of the first computing device, the proof of publication to the TEE; receiving, with at least one processor of the first computing device, the one-time key from the TEE; and determining, with at least one processor of the first computing device, the computation by decrypting the encrypted output with the one-time key.

[0021] In some non-limiting embodiments or aspects, the first portion of the one-time key that is published to the ledger may be encrypted with a public key associated with the second computing device.

[0022] In some non-limiting embodiments or aspects, the one-time key may be generated by the TEE from the first portion of the one-time key, the second portion of the one-time key, and a unique nonce associated with an instance of the computation.

[0023] In some non-limiting embodiments or aspects, the first portion of the one-time key may be generated from a private key associated with the first computing

device, and the second portion of the one-time key may be generated from a private key associated with the second computing device.

[0024] In some non-limiting embodiments or aspects, the first portion of the one-time key and the second portion of the one-time key may be communicated to the TEE in a first time period and may be configured to generate a plurality of one-time keys over successive instances of multi-party computation in subsequent time periods.

[0025] According to non-limiting embodiments or aspects, provided is a computer-implemented method for multi-party computation between more than two parties. The method includes communicating, with at least one processor of a first computing device, a first input to each of a plurality of TEEs. The method further includes receiving, with at least one processor of the first computing device from a TEE of the plurality of TEEs, an encrypted output of a computation completed by the TEE of the first input and a plurality of other inputs communicated to the TEE by a plurality of other party computing devices. A number of the plurality of other party computing devices is greater than a number of the plurality of TEEs. Each of a subset of the plurality of other party computing devices is associated with a TEE of the plurality of TEEs. The method further includes broadcasting, with at least one processor of the first computing device, and in response to receiving the encrypted output, a first party signature to each of the plurality of other party computing devices that is associated with a TEE of the plurality of TEEs. The method further includes receiving, with at least one processor of the first computing device, another party signature from each of the plurality of other party computing devices in response to the plurality of other party computing devices receiving the encrypted output. The method further includes communicating, with at least one processor of the first computing device, the first input to each of the plurality of other party computing devices that is not associated with a TEE of the plurality of TEEs. The method further includes, in response to receiving the plurality of other inputs from another party computing device of the plurality of other party computing devices: generating, with at least one processor of the first computing device, the computation of the first input and the plurality of other inputs. The method further includes, in response to not receiving the plurality of other inputs from any other party computing device of the plurality of other party computing devices: executing, with at least one processor of the first computing device, a fallback computation process using a TEE of the plurality of TEEs and a shared ledger to retrieve the computation of the first input and the plurality of other inputs.

[0026] According to non-limiting embodiments or aspects, provided is a system for multi-party computation between more than two parties. The system includes a first computing device and a trusted execution environment (TEE) associated with the first computing device. At least one processor of the first computing device is programmed or configured to execute one or more steps of the above-described method for multi-party computation between more than two parties.

[0027] According to non-limiting embodiments or aspects, provided is a computer program product for multi-party computation between more than two parties. The computer program product is stored on a first computing device and includes at least one non-transitory computer-readable medium including program instructions. The program instructions, when executed by at least one processor of the first computing device, cause the at least one processor to execute one or more steps of the above-described method for multi-party computation between more than two parties.

[0028] Other non-limiting embodiments or aspects of the present disclosure will be set forth in the following numbered clauses:

[0029] Clause 1: A system for multi-party computation between a first party and a second party, the system comprising: a first computing device associated with the first party; and a trusted execution environment (TEE) associated with the first computing device, wherein the first computing device is programmed or configured to: communicate, to the TEE, a first computation input and a first portion of a one-time key; receive, from the TEE, an encrypted output of a computation based on the first computation input and a second computation input, wherein the second computation input and a second portion of the one-time key are communicated to the TEE by a second computing device associated with the second party, and wherein the encrypted output is encrypted by the TEE with the one-time key; communicate the encrypted output to the second computing device; receive a digital signature from the second computing device indicating that the second computing device received the encrypted output; communicate the first portion of the one-time key to the second computing device; and (i) in response to receiving the second portion of the one-time key from the second computing device, determine the one-time key from the first portion of the one-time key and the second portion of the one-time key, and determine the computation by decrypting the encrypted output with the one-time key; or (ii) in response to not receiving the second portion of the one-time key from the second computing device, execute a fallback computation process using the TEE and a ledger

accessible to both the first computing device and the second computing device to determine the computation.

[0030] Clause 2: The system of clause 1, wherein the first computing device is further programmed or configured to, in the fallback computation process: publish, to the ledger, the first portion of the one-time key and the digital signature received from the second computing device; receive a proof of publication from the ledger; input the proof of publication to the TEE; receive the one-time key from the TEE; and determine the computation by decrypting the encrypted output with the one-time key.

[0031] Clause 3: The system of clause 1 or clause 2, wherein the first portion of the one-time key that is published to the ledger by the first computing device is encrypted with a public key associated with the second computing device.

[0032] Clause 4: The system of any of clauses 1-3, wherein the one-time key is generated by the TEE from the first portion of the one-time key, the second portion of the one-time key, and a unique nonce associated with an instance of the computation.

[0033] Clause 5: The system of any of clauses 1-4, wherein the first portion of the one-time key is generated from a private key associated with the first computing device, and wherein the second portion of the one-time key is generated from a private key associated with the second computing device.

[0034] Clause 6: The system of any of clauses 1-5, wherein the first portion of the one-time key and the second portion of the one-time key are communicated to the TEE in a first time period and are configured to generate a plurality of one-time keys over successive instances of multi-party computation in subsequent time periods.

[0035] Clause 7: The system of any of clauses 1-6, wherein the ledger comprises an append-only blockchain and is configured to be accessible only by computing devices that contribute to the computation.

[0036] Clause 8: A computer program product for multi-party computation between a first party and a second party, the computer program product being stored on a first computing device associated with the first party and comprising at least one non-transitory computer-readable medium including program instructions that, when executed by at least one processor of the first computing device, cause the at least one processor to: communicate, to a trusted execution environment (TEE) associated with the first computing device, a first computation input and a first portion of a one-time key; receive, from the TEE, an encrypted output of a computation based on the first computation input and a second computation input, wherein the second

computation input and a second portion of the one-time key are communicated to the TEE by a second computing device associated with the second party, and wherein the encrypted output is encrypted by the TEE with the one-time key; communicate the encrypted output to the second computing device; receive a digital signature from the second computing device indicating that the second computing device received the encrypted output; communicate the first portion of the one-time key to the second computing device; and (i) in response to receiving the second portion of the one-time key from the second computing device, determine the one-time key from the first portion of the one-time key and the second portion of the one-time key, and determine the computation by decrypting the encrypted output with the one-time key; or (ii) in response to not receiving the second portion of the one-time key from the second computing device, execute a fallback computation process using the TEE and a ledger accessible to both the first computing device and the second computing device to determine the computation.

[0037] Clause 9: The computer program product of clause 8, wherein the program instructions further cause the at least one processor to, in the fallback computation process: publish, to the ledger, the first portion of the one-time key and the digital signature received from the second computing device; receive a proof of publication from the ledger; input the proof of publication to the TEE; receive the one-time key from the TEE; and determine the computation by decrypting the encrypted output with the one-time key.

[0038] Clause 10: The computer program product of clause 8 or clause 9, wherein the first portion of the one-time key that is published to the ledger by the at least one processor is encrypted with a public key associated with the second computing device.

[0039] Clause 11: The computer program product of any of clauses 8-10, wherein the one-time key is generated by the TEE from the first portion of the one-time key, the second portion of the one-time key, and a unique nonce associated with an instance of the computation.

[0040] Clause 12: The computer program product of any of clauses 8-11, wherein the first portion of the one-time key is generated from a private key associated with the first computing device, and wherein the second portion of the one-time key is generated from a private key associated with the second computing device.

[0041] Clause 13: The computer program product of any of clauses 8-12, wherein the first portion of the one-time key and the second portion of the one-time key are

communicated to the TEE in a first time period and are configured to generate a plurality of one-time keys over successive instances of multi-party computation in subsequent time periods.

[0042] Clause 14: The computer program product of any of clauses 8-13, wherein the ledger comprises an append-only blockchain and is configured to be accessible only by computing devices that contribute to the computation.

[0043] Clause 15: A computer-implemented method for multi-party computation between a first party and a second party, the method comprising: communicating, with at least one processor of a first computing device associated with the first party, to a trusted execution environment (TEE) associated with the first computing device, a first computation input and a first portion of a one-time key; receiving, with at least one processor of the first computing device, from the TEE, an encrypted output of a computation based on the first computation input and a second computation input, wherein the second computation input and a second portion of the one-time key are communicated to the TEE by a second computing device associated with the second party, and wherein the encrypted output is encrypted by the TEE with the one-time key; communicating, with at least one processor of the first computing device, the encrypted output to the second computing device; receiving, with at least one processor of the first computing device, a digital signature from the second computing device indicating that the second computing device received the encrypted output; communicating, with at least one processor of the first computing device, the first portion of the one-time key to the second computing device; and in response to not receiving the second portion of the one-time key from the second computing device, executing, with at least one processor of the first computing device, a fallback computation process using the TEE and a ledger accessible to both the first computing device and the second computing device to determine the computation.

[0044] Clause 16: The method of clause 15, wherein the fallback computation process comprises: publishing, with at least one processor of the first computing device, to the ledger, the first portion of the one-time key and the digital signature received from the second computing device; receiving, with at least one processor of the first computing device, a proof of publication from the ledger; inputting, with at least one processor of the first computing device, the proof of publication to the TEE; receiving, with at least one processor of the first computing device, the one-time key

from the TEE; and determining, with at least one processor of the first computing device, the computation by decrypting the encrypted output with the one-time key.

[0045] Clause 17: The method of clause 15 or clause 16, wherein the first portion of the one-time key that is published to the ledger is encrypted with a public key associated with the second computing device.

[0046] Clause 18: The method of any of clauses 15-17, wherein the one-time key is generated by the TEE from the first portion of the one-time key, the second portion of the one-time key, and a unique nonce associated with an instance of the computation.

[0047] Clause 19: The method of any of clauses 15-18, wherein the first portion of the one-time key is generated from a private key associated with the first computing device, and wherein the second portion of the one-time key is generated from a private key associated with the second computing device.

[0048] Clause 20: The method of any of clauses 15-19, wherein the first portion of the one-time key and the second portion of the one-time key are communicated to the TEE in a first time period and are configured to generate a plurality of one-time keys over successive instances of multi-party computation in subsequent time periods.

[0049] Clause 21: A computer-implemented method for multi-party computation between more than two parties, the method comprising: communicating, with at least one processor of a first computing device, a first input to each of a plurality of TEEs; receiving, with at least one processor of the first computing device from a TEE of the plurality of TEEs, an encrypted output of a computation completed by the TEE of the first input and a plurality of other inputs communicated to the TEE by a plurality of other party computing devices, wherein a number of the plurality of other party computing devices is greater than a number of the plurality of TEEs, and wherein each of a subset of the plurality of other party computing devices is associated with a TEE of the plurality of TEEs; broadcasting, with at least one processor of the first computing device and in response to receiving the encrypted output, a first party signature to each of the plurality of other party computing devices that is associated with a TEE of the plurality of TEEs; receiving, with at least one processor of the first computing device, another party signature from each of the plurality of other party computing devices in response to the plurality of other party computing devices receiving the encrypted output; communicating, with at least one processor of the first computing device, the first input to each of the plurality of other party computing devices that is

not associated with a TEE of the plurality of TEEs; and, in response to receiving the plurality of other inputs from another party computing device of the plurality of other party computing devices: generating, with at least one processor of the first computing device, the computation of the first input and the plurality of other inputs; or, in response to not receiving the plurality of other inputs from any other party computing device of the plurality of other party computing devices: executing, with at least one processor of the first computing device, a fallback computation process using a TEE of the plurality of TEEs and a shared ledger to retrieve the computation of the first input and plurality of other inputs.

[0050] Clause 22: A system for multi-party computation between more than two parties, the system comprising a first computing device and a trusted execution environment (TEE) associated with the first computing device, wherein at least one processor of the first computing device is programmed or configured to execute one or more steps of clause 21.

[0051] Clause 23: A computer program product for multi-party computation between more than two parties, the computer program product being stored on a first computing device and comprising at least one non-transitory computer-readable medium including program instructions, wherein the program instructions, when executed by at least one processor of the first computing device, cause the at least one processor to execute one or more steps of clause 21.

[0052] These and other features and characteristics of the present disclosure, as well as the methods of operation and functions of the related elements of structures and the combination of parts and economies of manufacture, will become more apparent upon consideration of the following description and the appended claims with reference to the accompanying drawings, all of which form a part of this specification, wherein like reference numerals designate corresponding parts in the various figures. It is to be expressly understood, however, that the drawings are for the purpose of illustration and description only and are not intended as a definition of the limits of the present disclosure. As used in the specification and the claims, the singular form of "a," "an," and "the" include plural referents unless the context clearly dictates otherwise.

BRIEF DESCRIPTION OF THE DRAWINGS

[0053] Additional advantages and details of the disclosure are explained in greater detail below with reference to the exemplary embodiments that are illustrated in the accompanying figures, in which:

[0054] FIG. 1 is a diagram of a non-limiting embodiment or aspect of an environment in which systems, apparatuses, and/or methods, as described herein, may be implemented;

[0055] FIG. 2 is a diagram of a non-limiting embodiment or aspect of components of one or more devices of FIG. 1;

[0056] FIG. 3 is a flow diagram of a non-limiting embodiment or aspect of a method for secure real-time n-party computation;

[0057] FIG. 4 is a flow diagram of a non-limiting embodiment or aspect of a method for secure real-time n-party computation;

[0058] FIG. 5 is a schematic diagram of a non-limiting embodiment or aspect of a system for secure real-time n-party computation;

[0059] FIG. 6 is a schematic diagram of a non-limiting embodiment or aspect of a system for secure real-time n-party computation;

[0060] FIG. 7 is a set of pseudocode describing a method for secure real-time n-party computation, according to non-limiting embodiments or aspects;

[0061] FIG. 8 is a set of pseudocode describing a method for secure real-time n-party computation, according to non-limiting embodiments or aspects;

[0062] FIG. 9 is a set of pseudocode describing a method for secure real-time n-party computation, according to non-limiting embodiments or aspects;

[0063] FIG. 10 is a set of pseudocode describing a method for secure real-time n-party computation, according to non-limiting embodiments or aspects;

[0064] FIG. 11 is a set of pseudocode describing a method for secure real-time n-party computation, according to non-limiting embodiments or aspects;

[0065] FIG. 12 is a set of pseudocode describing a method for secure real-time n-party computation, according to non-limiting embodiments or aspects;

[0066] FIG. 13 is a set of pseudocode describing a method for secure real-time n-party computation, according to non-limiting embodiments or aspects;

[0067] FIG. 14 is a set of pseudocode describing a method for secure real-time n-party computation, according to non-limiting embodiments or aspects;

[0068] FIG. 15 is a set of pseudocode describing a method for secure real-time n-party computation, according to non-limiting embodiments or aspects;

[0069] FIG. 16 is a set of pseudocode describing a method for secure real-time n-party computation, according to non-limiting embodiments or aspects; and

[0070] FIG. 17 is a set of pseudocode describing a method for secure real-time n-party computation, according to non-limiting embodiments or aspects.

[0071] It should be appreciated by those skilled in the art that any block diagrams herein represent conceptual views of illustrative systems embodying the principles of the present subject matter. Similarly, it may be appreciated that any flow charts, flow diagrams, state transition diagrams, pseudo code, and the like represent various processes which may be substantially represented in computer readable medium and executed by a computer or processor, whether or not such computer or processor is explicitly shown.

DETAILED DESCRIPTION

[0072] For purposes of the description hereinafter, the terms "upper", "lower", "right", "left", "vertical", "horizontal", "top", "bottom", "lateral", "longitudinal," and derivatives thereof shall relate to non-limiting embodiments or aspects as they are oriented in the drawing figures. However, it is to be understood that non-limiting embodiments or aspects may assume various alternative variations and step sequences, except where expressly specified to the contrary. It is also to be understood that the specific devices and processes illustrated in the attached drawings, and described in the following specification, are simply exemplary embodiments or aspects. Hence, specific dimensions and other physical characteristics related to the embodiments or aspects disclosed herein are not to be considered as limiting.

[0073] No aspect, component, element, structure, act, step, function, instruction, and/or the like used herein should be construed as critical or essential unless explicitly described as such. Also, as used herein, the articles "a" and "an" are intended to include one or more items and may be used interchangeably with "one or more" and "at least one." Furthermore, as used herein, the term "set" is intended to include one or more items (e.g., related items, unrelated items, a combination of related and unrelated items, etc.) and may be used interchangeably with "one or more" or "at least one." Where only one item is intended, the term "one" or similar language is used.

Also, as used herein, the terms "has," "have," "having," or the like are intended to be open-ended terms. Further, the phrase "based on" is intended to mean "based at least partially on" unless explicitly stated otherwise.

[0074] Some non-limiting embodiments or aspects are described herein in connection with thresholds. As used herein, satisfying a threshold may refer to a value being greater than the threshold, more than the threshold, higher than the threshold, greater than or equal to the threshold, less than the threshold, fewer than the threshold, lower than the threshold, less than or equal to the threshold, equal to the threshold, and/or the like.

[0075] As used herein, the term "acquirer institution" may refer to an entity licensed and/or approved by a transaction service provider to originate transactions (e.g., payment transactions) using a payment device associated with the transaction service provider. The transactions the acquirer institution may originate may include payment transactions (e.g., purchases, original credit transactions (OCTs), account funding transactions (AFTs), and/or the like). In some non-limiting embodiments or aspects, an acquirer institution may be a financial institution, such as a bank. As used herein, the term "acquirer system" may refer to one or more computing devices operated by or on behalf of an acquirer institution, such as a server computer executing one or more software applications.

[0076] As used herein, the term "account identifier" may include one or more primary account numbers (PANs), tokens, or other identifiers associated with a customer account. The term "token" may refer to an identifier that is used as a substitute or replacement identifier for an original account identifier, such as a PAN. Account identifiers may be alphanumeric or any combination of characters and/or symbols. Tokens may be associated with a PAN or other original account identifier in one or more data structures (e.g., one or more databases, and/or the like) such that they may be used to conduct a transaction without directly using the original account identifier. In some examples, an original account identifier, such as a PAN, may be associated with a plurality of tokens for different individuals or purposes.

[0077] As used herein, the term "communication" may refer to the reception, receipt, transmission, transfer, provision, and/or the like, of data (e.g., information, signals, messages, instructions, commands, and/or the like). For one unit (e.g., a device, a system, a component of a device or system, combinations thereof, and/or the like) to be in communication with another unit means that the one unit is able to directly or

indirectly receive information from and/or transmit information to the other unit. This may refer to a direct or indirect connection (e.g., a direct communication connection, an indirect communication connection, and/or the like) that is wired and/or wireless in nature. Additionally, two units may be in communication with each other even though the information transmitted may be modified, processed, relayed, and/or routed between the first and second unit. For example, a first unit may be in communication with a second unit even though the first unit passively receives information and does not actively transmit information to the second unit. As another example, a first unit may be in communication with a second unit if at least one intermediary unit processes information received from the first unit and communicates the processed information to the second unit.

[0078] As used herein, the term “computing device” may refer to one or more electronic devices configured to process data. A computing device may, in some examples, include the necessary components to receive, process, and output data, such as a processor, a display, a memory, an input device, a network interface, and/or the like. A computing device may be a mobile device. As an example, a mobile device may include a cellular phone (e.g., a smartphone or standard cellular phone), a portable computer, a wearable device (e.g., watches, glasses, lenses, clothing, and/or the like), a personal digital assistant (PDA), and/or other like devices. A computing device may also be a desktop computer or other form of non-mobile computer. An “application” or “application program interface” (API) may refer to computer code or other data stored on a computer-readable medium that may be executed by a processor to facilitate the interaction between software components, such as a client-side front-end and/or server-side back-end for receiving data from the client. An “interface” may refer to a generated display, such as one or more graphical user interfaces (GUIs) with which a user may interact, either directly or indirectly (e.g., through a keyboard, mouse, etc.).

[0079] As used herein, the terms “electronic wallet” and “electronic wallet application” refer to one or more electronic devices and/or software applications configured to initiate and/or conduct payment transactions. For example, an electronic wallet may include a mobile device executing an electronic wallet application, and may further include server-side software and/or databases for maintaining and providing transaction data to the mobile device. An “electronic wallet provider” may include an entity that provides and/or maintains an electronic wallet for a customer, such as

Google Pay®, Android Pay®, Apple Pay®, Samsung Pay®, and/or other like electronic payment systems. In some non-limiting examples, an issuer bank may be an electronic wallet provider.

[0080] As used herein, the term "issuer institution" may refer to one or more entities, such as a bank, that provide accounts to customers for conducting transactions (e.g., payment transactions), such as initiating credit and/or debit payments. For example, an issuer institution may provide an account identifier, such as a primary account number (PAN), to a customer that uniquely identifies one or more accounts associated with that customer. The account identifier may be embodied on a portable financial device, such as a physical financial instrument, e.g., a payment card, and/or may be electronic and used for electronic payments. The term "issuer system" refers to one or more computer devices operated by or on behalf of an issuer institution, such as a server computer executing one or more software applications. For example, an issuer system may include one or more authorization servers for authorizing a transaction.

[0081] As used herein, the term "merchant" may refer to an individual or entity that provides goods and/or services, or access to goods and/or services, to customers based on a transaction, such as a payment transaction. The term "merchant" or "merchant system" may also refer to one or more computer systems operated by or on behalf of a merchant, such as a server computer executing one or more software applications. A "point-of-sale (POS) system," as used herein, may refer to one or more computers and/or peripheral devices used by a merchant to engage in payment transactions with customers, including one or more card readers, scanning devices (e.g., code scanners), Bluetooth® communication receivers, near-field communication (NFC) receivers, radio frequency identification (RFID) receivers, and/or other contactless transceivers or receivers, contact-based receivers, payment terminals, computers, servers, input devices, and/or other like devices that can be used to initiate a payment transaction.

[0082] As used herein, the term "payment device" may refer to a payment card (e.g., a credit or debit card), a gift card, a smartcard, smart media, a payroll card, a healthcare card, a wristband, a machine-readable medium containing account information, a keychain device or fob, an RFID transponder, a retailer discount or loyalty card, a cellular phone, an electronic wallet mobile application, a PDA, a pager, a security card, a computing device, an access card, a wireless terminal, a transponder, and/or the like. In some non-limiting embodiments or aspects, the

payment device may include volatile or non-volatile memory to store information (*e.g.*, an account identifier, a name of the account holder, and/or the like).

[0083] As used herein, the term “payment gateway” may refer to an entity and/or a payment processing system operated by or on behalf of such an entity (*e.g.*, a merchant service provider, a payment service provider, a payment facilitator, a payment facilitator that contracts with an acquirer, a payment aggregator, and/or the like), which provides payment services (*e.g.*, transaction service provider payment services, payment processing services, and/or the like) to one or more merchants. The payment services may be associated with the use of portable financial devices managed by a transaction service provider. As used herein, the term “payment gateway system” may refer to one or more computer systems, computer devices, servers, groups of servers, and/or the like, operated by or on behalf of a payment gateway.

[0084] As used herein, the term “server” may refer to or include one or more computing devices that are operated by or facilitate communication and processing for multiple parties in a network environment, such as the Internet, although it will be appreciated that communication may be facilitated over one or more public or private network environments and that various other arrangements are possible. Further, multiple computing devices (*e.g.*, servers, POS devices, mobile devices, etc.) directly or indirectly communicating in the network environment may constitute a “system.” Reference to “a server” or “a processor,” as used herein, may refer to a previously-recited server and/or processor that is recited as performing a previous step or function, a different server and/or processor, and/or a combination of servers and/or processors. For example, as used in the specification and the claims, a first server and/or a first processor that is recited as performing a first step or function may refer to the same or different server and/or a processor recited as performing a second step or function.

[0085] As used herein, the term “transaction service provider” may refer to an entity that receives transaction authorization requests from merchants or other entities and provides guarantees of payment, in some cases through an agreement between the transaction service provider and an issuer institution. For example, a transaction service provider may include a payment network such as Visa® or any other entity that processes transactions. The term “transaction processing system” may refer to one or more computer systems operated by or on behalf of a transaction service

provider, such as a transaction processing server executing one or more software applications. A transaction processing server may include one or more processors and, in some non-limiting embodiments or aspects, may be operated by or on behalf of a transaction service provider.

[0086] As used herein, an electronic payment processing network may refer to the communications between one or more entities for processing the transfer of monetary funds to one or more transactions. The electronic payment processing network may include a merchant system, an acquirer system, a transaction service provider, and an issuer system.

[0087] Non-limiting embodiments or aspects of the present disclosure are directed to a system, method, and computer program product for secure real-time n -party computation. As used herein, "computation" may refer to any joint effort between computing devices where independent inputs are combined to produce a result based on each input. For example, computation may include, but is not limited to, mathematical computation, data compilation and exchange, execution of contracts, transactional exchange, and/or the like. It will be appreciated that described systems and methods for two-party exchange may be expanded to n -party computations, where $n > 1$. Described embodiments and aspects improve upon prior art technical solutions by providing an architecture for secure, fairness-assured joint computation. The described systems and methods provide fair secure computation with fallback mechanisms if a participating party should cease cooperation. No third party intermediary computing devices need to be involved to assure fairness, reducing communication time associated with outside transmissions, and decreasing the number of party resources involved. Furthermore, the described systems and methods remove the possibility of a third party intermediary computing device colluding with one of the parties, thereby threatening security and fairness. In the present described methods, a corrupt party cannot receive the final output without the assurance that other participating parties also have access to joint computation. Moreover, the computation and exchange of information can happen in real-time to other processes (e.g., transaction processing) by not relying on third party verification services.

[0088] In addition to the above-described technical improvements, described systems and methods provide a more computationally efficient process for multi-party computation by not relying on a shared ledger (e.g., a blockchain), except and until

any party proves to be malicious and aborts early. This configuration reduces the number of computational resources and independent communications required to execute fair multi-party computation.

Protocol Terminology

[0089] As used herein, a “broadcast” protocol for n parties $\{P_1, \dots, P_n\}$, where a commander $P_c \in \{P_1, \dots, P_n\}$ inputs a message m , terminates with properties of agreement and validity. Agreement states that all honest parties output the same value, and validity states that all parties output m if the commander is honest. The term “broadcast” may refer to the Dolev-Strong broadcast protocol.

[0090] As used herein, a “trusted execution environment” (TEE) may refer to a secure area of a main processor. A TEE may guarantee that code and data loaded inside is protected with respect to confidentiality and integrity. An enclave program is an isolated region of memory, containing both code and data, protected by the TEE platform (where trust is only placed in the processor manufacturer). On TEE platforms (e.g., Intel® SGX, Sanctum™, etc.), the central processing unit (CPU) monitors all memory accesses to ensure that non-enclave software (e.g., operating system, Hypervisor®, basic input/output system (BIOS) software, system management mode (SMM) firmware, etc.) cannot access the enclave’s memory. In addition to isolated execution, as used herein, the TEE platform may provide a primitive for remote attestation. At any time, the enclave software may request a signed message, called a “quote”, binding an enclave-supplied value to that enclave’s code identity (e.g., as a hash-based measurement).

[0091] As used herein, described protocols may make use of a shared, append-only ledger. This ledger may be referred to in the context of a “bulletin board” abstraction. The shared ledger allows parties to get the contents of the ledger and post arbitrary strings on it. Furthermore, on successfully publishing a string on the shared ledger, any party can request a proof (e.g., a publicly verifiable proof) that the string was indeed published, and the shared ledger guarantees that the string will never be modified or deleted. Hence, the shared ledger may also be referred to as an append-only ledger. In some non-limiting embodiments or aspects, the shared ledger may be instantiated using fork-less blockchains, e.g., permissioned blockchains, or by blockchains based on proof-of-stake.

[0092] In some non-limiting embodiments or aspects, described protocols may make use of an environment where the parties have a fully connected network with point-to-point communication channels and can communicate in synchronous rounds (*e.g.*, stages of steps executed sequentially over time). Moreover, described protocols may include a public key infrastructure amongst the parties, where each part has a public key pk_i and a private key sk_i . The honest parties use a key generation algorithm to generate the key-pair. Malicious parties may generate arbitrary keys. Prior to running the coin-tossing protocol (described further below), all parties have the same vector of public keys of all other participants. The coin-tossing protocol may include a public key encryption scheme and digital signature scheme, which enables the parties to communicate via authenticated channels, thereby preventing an adversary from modifying messages or injecting messages that claim to be from a different (honest) party.

Two-Party Protocol

[0093] In some non-limiting embodiments or aspects, described is a protocol tailored to the two-party setting, which may be simpler and more round-efficient (*e.g.*, fewer stages of steps) compared to the generalized n -party protocol. The goal of the two-party protocol is as follows. The parties wish to perform a joint secure computation of some function f over private inputs from both parties, with fairness: either both parties get the output of f , or no one does. The described two-party protocol improves over prior techniques, in that the protocol avoids posting to the shared ledger in the case where the parties behave honestly and follow the protocol's steps (thereby reducing the recruitment of additional computer resources to the multi-party computation). Second, the described two-party protocol enables pre-processing, where subsequent evaluations of f (over different inputs from the two parties) can amortize the work from the earlier phases of the protocol—this helps to reduce the round complexity (*e.g.*, fewer steps per stage) of the protocol.

[0094] The described two-party protocol reduces the problem of fair secure computation to fair reconstruction of an additive secret sharing scheme, by encrypting the output of f under a randomly chosen key k , and distributing *xor* shares of k to both parties—*xor* ensures that each share reveals zero information about k . In that sense, the reconstruction protocol's responsibility is to ensure fair exchange of the shares held by each party. It will be appreciated that two-party fair exchange implies two-

party fair computation, and vice versa. In this approach, both parties are on a shared ledger, and one of the parties has access to a TEE.

[0095] The protocol includes a fallback mechanism (to be invoked by the TEE-possessing party) containing a TEE-ledger interaction that gives the decryption key k to all parties—during malicious behavior, the non-TEE party may wait for the other party to perform this TEE-ledger interaction, or else, the non-TEE party is guaranteed that no party gets k , thus ensuring fairness to both parties. The described protocol has the parties share their keys with each other in a particular sequence, where regardless of when the other party deviates from the protocol (*e.g.*, aborts), the other party has a mechanism to terminate the protocol with fairness.

[0096] The protocol commences with a pre-processing step which loads both parties' computational shares (*e.g.*, inputs) in the TEE, which is useful during the TEE-ledger interaction in the event of malicious behavior. This step may only need to be performed once for multiple invocations of the two-party protocol (*e.g.*, for multiple evaluations of the function f). The pre-processing phase "loads" the TEE with a secret key from which k is derived for each instance of the two-party protocol. To that end, instead of choosing k randomly, k may be derived from long-term secret keys held by both parties and a unique nonce for that instance of the two-party protocol. Then, the (unfair) MPC may receive the keys, *e.g.*, msk_1 and msk_2 , as input from both parties, along with the common value of nonce, and may produce the output encrypted under:

Formula 1

$$k \triangleq \text{PRF}(msk_1 \oplus msk_2, \text{nonce})$$

[0097] The two-party protocol terminates in both parties getting the decryption key, which is no use if the parties are not given the encrypted outputs. Therefore, as a first step in the fair exchange protocol, the second party P_2 sends a signature signifying P_2 's receipt of the encrypted output. Should P_2 not send this message, the protocol may abort and no party may learn the output. Although P_2 can abort if P_2 does not attain the encrypted output from the two-party protocol, a malicious first party P_1 may still carry out the fallback mechanism, which uses a TEE-ledger interaction to provide the decryption key to both parties with fairness, and which would not be useful to P_2 without the encrypted output. The signature protects P_2 against P_1 maliciously initiating the fallback TEE-ledger interaction. It will be appreciated that it suffices for

only P_1 to receive P_2 's signature, since only P_1 may use P_1 's TEE to trigger the TEE-ledger interaction.

[0098] The two-party protocol further imposes an order in which the parties reveal their shares to each other. Since only P_1 has a TEE and can invoke the fallback mechanism as a recourse against malicious behavior, P_1 is required to first send P_1 's share to P_2 . It will be appreciated that alternative approaches, such as P_2 sharing first or simultaneous sharing of the shares, opens the possibility of P_1 attaining all shares, leaving P_2 without any recourse.

[0099] If P_2 receives a share from P_1 , P_2 sends P_2 's share to P_1 . While it is possible that P_1 could abort the protocol (in which case, no party gets both shares), there is also a possibility that P_1 may maliciously avoid sending the messages to P_2 and may instead trigger the fallback TEE-ledger mechanism. For that reason, if P_2 does not receive any message, P_2 reads the shared ledger thereafter, searching for messages from a potential trigger of the fallback mechanism. If either party deviates from the fair exchange protocol, the fallback mechanism may be employed. P_1 is required to post a message to the shared ledger containing the aforementioned signature from P_2 , and the decryption key (encrypted under P_2 's public key). Both parties may retrieve the ledger entry to complete the protocol. For P_2 , the ledger entry contains the decryption key, which P_2 outputs. Meanwhile, P_1 provides the proof of publication (denoting that the entry was successfully posted on the shared ledger) to the local TEE, which proceeds to verify the proof and release the decryption key locally to P_1 .

[0100] For the purpose of further illustration, two-party fair reconstruction, including a first party (P_1) and a second party (P_2), may be further described as a process occurring in a series of rounds. Rounds, as used herein, may refer to sequential stages in which one or more steps of the described method occur. For two-party fair reconstruction, P_1 may launch a TEE, which attests itself to P_2 and establishes a session key sk with P_2 . P_2 may use this session key sk to encrypt its inputs. The TEE may also generate a master session key msk , which is always kept within the TEE and is used to derive the one-time key $k = PRF(msk, id)$, where id is the function evaluation's unique identifier. For illustrative purposes, P_1 may be associated with a TEE, and P_2 may not be associated with a TEE. It may be assumed that both P_1 and P_2 have access to a shared ledger (L). The function "Enc" may denote an encryption function, the function "PRF" may denote a pseudorandom function, and the function "Sign" may denote a digital signature function.

[0101] In Round 0, P_1 waits for input $Enc(sk, x_2) || id$ from P_2 . Also in Round 0, P_2 sends input $Enc(sk, x_2) || id$ to P_1 . In Round 1, P_1 sends encrypted output $m1 = Enc(k, y)$ to P_2 , where $k = PRF(msk, id)$. Also in Round 1, P_2 waits for encrypted output $m1 = Enc(k, y)$ from P_1 . If P_2 does not receive $m1$, then P_2 may abort the reconstruction process. In Round 2, P_1 waits for signature $m2 = sign(sk_2, id)$ from P_2 . If no message is received, then P_1 may abort the reconstruction process. Also in Round 2, P_2 sends signature $m2 = sign(sk_2, id)$ to P_1 signifying that P_2 got the encrypted output from Round 1. In Round 3, P_1 sends $m3 = Enc(pk_2, k)$ to P_2 . Also in Round 3, P_2 waits for $m3 = Enc(pk_2, k)$ from P_1 . In Round 4, P_1 waits for $m4 = Enc(pk_1, k)$ from P_2 . Also in Round 4, if P_2 received $m3$ from P_1 , then P_2 may send $m4 = Enc(pk_1, k)$ to P_1 and return k . If P_2 did not receive $m3$, P_2 may sleep (e.g., do nothing). In Round 5, if P_1 does not receive $m4$, then P_1 may trigger a fallback process by posting $Enc(pk_2, k)$ to L . Also in Round 5, P_2 may sleep. In Round 6, P_1 receives proof of publication from L and provides the proof of publication to TEE to get k . Also in Round 6, party P_2 receives proof of publication and entry from L to get k .

N-Party Fair Computation

[0102] The present disclosure further provides a generalized protocol for n -party fair computation. The goal of the protocol is to perform a joint secure computation of some function f over private inputs from all n parties, with fairness. A fair exchange indicates that either all the parties receive the output of function f , or no party does.

[0103] The described n -party protocol is designed for an optimistic execution off-chain, where the protocol avoids posting to the shared ledger when the parties behave honestly and follow the protocol's steps. This reduces computational resources when compared to other methods that require posts to the shared ledger throughout the process. Second, the protocol enables pre-processing, where subsequent evaluations of f (over different inputs from the n parties) can amortize the work from the earlier phases of the protocol. This helps to improve the round complexity of the protocol.

[0104] In the n -party protocol, the problem of fair secure computation is reduced to fair reconstruction of an additive secret sharing scheme, by encrypting the output of f under a random key k , and distributing xor shares $k_1, \dots, k_n$ of k to all n parties, such that:

Formula 2

$$k = k_1 \oplus \dots \oplus k_n$$

Distributing xor shares ensures that each share reveals zero information about k . In that sense, the reconstruction protocol's responsibility is to ensure fair exchange of the shares held by each party. The n -party protocol may give each party the encrypted output and that party's share of k . Moreover, all parties are on a shared ledger, and t out of n parties have access to a TEE (for a corruption threshold t in a static corruption model).

[0105] The n -party protocol relies on a fallback mechanism, to be invoked by any TEE-possessing party that detects malicious behavior, containing a TEE-ledger interaction that gives the decryption key k to all parties. During malicious behavior, the non-TEE parties may wait for some honest TEE-possessing party to perform this TEE-ledger interaction, or else, they are guaranteed that no party gets k , thus ensuring fairness to all n parties. The technical challenge addressed by this protocol's design is to have the parties share their keys with each other in a safe sequence, where regardless of when a malicious party deviates from the protocol (e.g., aborts after receiving sufficient shares), the other parties have a mechanism to carry out the remaining protocol steps with the guarantee of fairness.

[0106] The n -party protocol may be executed between n parties $P_1, \dots, P_n$. At a high-level, if all parties behave honestly (optimistic setting), the protocol terminates in a fixed number of rounds with no interaction with the shared ledger. Moreover, if any party behaves maliciously, the protocol either terminates (when there is a guarantee that no party could have recovered the key k), or uses a TEE-ledger interaction to deliver k to all parties with fairness.

[0107] Similar to the two-party protocol, the n -party protocol may be commenced with a pre-processing step where parties submit their secrets to all t TEEs, which is useful during the fallback TEE-ledger mechanism in the event of malicious behavior. Similar to the two-party protocol, the output of f is encrypted under a key k that is not chosen randomly, but is instead generated from long-term secret keys held by all parties. Specifically, the (unfair) multi-party computation protocol encrypts the output of f under key:

Formula 3

$$k \doteq \text{PRF}(\text{msk}_1 \oplus \dots \oplus \text{msk}_n, \text{nonce})$$

where each party P_i maintains a long term secret key msk_i and inputs said long-term secret key to each instance of the protocol. Hence, the pre-processing step is amortized across several instances of the unfair multi-party computation protocol and the fair reconstruction protocols.

[0108] The n -party protocol guarantees that all n parties (or no parties) attain the decryption key, but that is of no use if any subset of the parties do not get the encrypted output from the (unfair) multi-party computation protocol execution. Therefore, as a first step in the fair reconstruction protocol, each party sends a signature signifying their receipt of the encrypted output. The signature is sent only to the t parties with TEEs, and having all n signatures is a prerequisite for the fallback TEE-ledger mechanism that gives k to all parties. It will be appreciated that a malicious party can later carry out the fallback mechanism that uses a TEE-ledger interaction to provide the decryption key to all parties with fairness, but that mechanism is of no use to a party that does not receive the encrypted output. Therefore, the signature protects non-malicious parties from malicious parties initiating the fallback TEE-ledger interaction.

[0109] The n -party protocol imposes an order in which the parties reveal the shares to each other. Since only parties with TEEs can invoke the fallback mechanism (described below) as a recourse against malicious behavior, the n -party protocol imposes an order where each party $P_i \in \{P_1, \dots, P_n\}$ first sends its share of k to each $P_j \in \{P_1, \dots, P_n\}$. Note that alternative approaches, such as P_i sharing first or simultaneous sharing of the shares, open the possibility of some malicious party P_i attaining all shares, leaving an honest party P_j without any recourse.

[0110] On receiving all t shares, each non-TEE party P_j may send its share of k to each TEE-possessing party P_i . If a P_j did not receive all t shares, then it may transition to the fallback mechanism (described below). While it is possible that P_j did not receive a share because a TEE party P_i aborted the protocol, it may be possible that P_i acted maliciously and withheld its share k_i from P_j and instead proceed to trigger the fallback TEE-ledger interaction. Therefore, party P_j cannot simply abort, but must read the shared ledger in the subsequent time period for a TEE-ledger interaction, which is designed to release k to all parties.

[0111] Any party to get all n shares may broadcast k to all parties. For example, if any honest party (TEE or non-TEE) receives all n shares of k , then said party broadcasts k to all parties. If the party did not get all n shares, then said party may attempt to receive k from an honest party. If the party still does not get k , then it may enter the fallback mechanism.

[0112] The fallback mechanism constitutes a TEE-ledger interaction. Any TEE-possessing party P_i can post a message to the shared ledger including all n signatures. The signatures not only indicate that all parties have the encrypted output, but posting them on the ledger enables any of the other TEE parties to carry out the next steps of the fallback mechanism. It will be appreciated that it is possible that a malicious party did not send its signature to a TEE party at the beginning of the n -party protocol. In the next steps of the fallback mechanism, reading the ledger entry provides a proof of publication that indicates that the entry was successfully posted on the shared ledger and is accessible to all parties. On providing this proof to the TEE, any TEE-possessing party can advance the protocol by posting a second message on the shared ledger. This message contains an encryption key k under the public keys of all non-TEE parties. Finally, all parties may retrieve the ledger entry to complete the fallback mechanism. For non-TEE parties, the ledger entry contains the decryption key. On the other hand, each TEE-possessing party provides the proof of publication (denoting that the entry was successfully posted on the shared ledger) to the local TEE, which proceeds to verify the proof and release the decryption key locally to that party.

N-Party Fair Coin Tossing

[0113] Further described is a fair n -party coin tossing protocol, for arbitrary n . The described coin tossing protocol lets each party P_i input a chosen (random) bit b_i , and the honest parties output the xor of bits from all honest parties. The coin tossing protocol is based on the premise that parties perform the protocol steps obliviously, *e.g.*, without any knowledge of the bits from the other parties, until the final step of the protocol returns the output. In other words, the attacker does not learn the inputs of honest parties from their messages until the last step where the uniformly random output is produced. This is achieved by having parties broadcast an encryption of their chosen bits under a key known only to the t enclaves, *e.g.*, even the parties who control the TEEs do not learn the bits. Next, any party with a TEE broadcasts the enclave

private key to all parties, allowing each party to decrypt the chosen bits of all other parties and compute the output.

[0114] The t parties with TEE nodes launch an enclave program on their local machine. As a first step, the coin-tossing protocol generates ephemeral public key encryption (PKE) keys within enclaves, and communicates the public key material to all parties for establishing a secure channel with the enclave. To that end, each enclave executes:

Formula 4

$$(epk_i, esk_i) \leftarrow \text{PKE.Keygen}(\lambda)$$

and returns enclave public key epk_i to the local party P_i while keeping enclave private key esk_i in enclave private memory. Then, all of the t parties engage in a broadcast protocol (in parallel) to send their epk_i to all parties $P_1, \dots, P_t$. Specifically, each $P_i \in \{P_1, \dots, P_t\}$ invokes the broadcast protocol with the following message:

Formula 5

$$p_i : \text{quote}_{\text{HW}}(\mathcal{E}_{\text{PKE}}, epk_i), \text{ where } (epk_i, esk_i) \leftarrow \text{PKE.Keygen}(1^\lambda)$$

[0115] Each instance of the broadcast protocol requires $t+1$ rounds (where t is the corruption threshold and also the number of parties with a TEE), and all t parties execute an instance of the protocol in parallel, where each party acts as the commander in their own instance of the protocol.

[0116] Since the later steps of the coin tossing protocol use broadcast of encrypted messages, the protocol establishes a shared key for use by all parties, whose private component will be kept within the enclaves (to prevent the parties from prematurely learning other parties' inputs). By establishing a shared key for all future encryptions by all parties, the protocol may progress even when some parties with TEEs abort. The protocol may also optimize the message lengths as they contain the encryption under only one key. To that end, the protocol establishes a PKE key that is shared with all enclaves and is used by all parties to encrypt their bits. The enclaves are tasked with holding onto this key until all parties have broadcasted their chosen bits, preventing any party from learning an honest party's chosen bit. One party is selected as a leader (e.g., P_1) to generate a PKE key from within its enclave and broadcast the public key material to all parties and the secret key material to the enclaves (by encrypting under their keys from the previous step). Should the leader act maliciously

(e.g., the honest parties output $\perp$ at the end of the broadcast protocol), the protocol retries with the next leader, where leaders iterate from P_1 to P_t . If all t parties with TEEs act maliciously during this step, then the remaining $n - t$ honest parties execute an honest variant of the fair coin tossing protocol (where the honest parties simply broadcast their chosen bits to other honest parties in the clear).

[0117] All n parties must broadcast their chosen bits to every other party, but they cannot simply expose these values in the clear lest the adversary bias the final output. Given that a shared public key pk has been established, each party P_i signs and encrypts its bit b_i under the key pk . The signature prevents copycat attacks, where an attacker copies an honest party's bit and biases the output. Each party P_i further broadcasts the signed ciphertext to the other parties:

Formula 6

$$p_i : \text{PKE.Enc}(pk, \text{S.Sig}(sk_{p_i}, b_i))$$

Ciphertexts can only be decrypted within the TEEs, as all copies of the private key sk are stored within enclave memory. Security is achieved (e.g., adversarial bias is prevented) by having the parties carry out the protocol while remaining oblivious of each other's chosen bits. Hence, a corrupt party may choose to deviate from the protocol but must make this decision without knowledge of the honest parties' inputs. This allows the final output of the honest parties' to be unbiased.

[0118] The broadcast protocol ensures that all honest parties have the encrypted bits of all other honest parties. To compute the final output bit, which is the xor of all parties' bits, the private key sk must be exposed out of the enclave and shared with all the parties. To that end, each party with an enclave extracts encryptions of sk under each of the non-enclave public keys from its local enclave and broadcasts this to all parties, as indicated below.

Formula 7

$$\text{quote}_{\text{HW}}(\mathcal{E}_{\text{ct}}, \{ \text{PKE.Enc}(pk_i, sk) \}_{i>t})$$

At the end of this phase, honest parties without a TEE agree on whether any of them obtained the vector of encryptions of sk . If none of them received the encryptions, then this means that all t parties with TEEs acted maliciously during this step. In this case, the remaining $n - t$ honest parties may execute an honest variant of the fair coin tossing protocol (e.g., one honest party picks a random bit and sends it to all other

parties). This works because, at this point in the protocol, the t malicious parties (e.g., all those possessing TEEs) are unaware of the final output, and thus their decision to abort is independent of the value of the coin. On the other hand, even when one party with a TEE acts honestly, then all honest parties without a TEE obtain sk from the broadcast and are able to decrypt the encryptions and obtain the output. In a scenario with at least one honest party with a TEE, said parties may then instruct their enclave to release sk following which they are able to decrypt the encryptions and obtain their output.

Illustrated Non-Limiting Embodiments and Aspects

[0119] Referring now to FIG. 1, illustrated is a diagram of an example environment 100 in which devices, systems, and/or methods, described herein, may be implemented. As shown in FIG. 1, environment 100 includes a first computing device 102 associated with a first party, a trusted execution environment (TEE) 104, one or more other computing devices 106 associated with parties other than the first party, a shared ledger 108, and a communication network 110. The TEE 104 may be associated with and/or implemented on the same device as the first computing device 102. One or more other ledgers 108 may be used, which may be accessible by one or more other computing devices 106. The shared ledger 108 may include an append-only blockchain, e.g., permissioned blockchains, and may be configured to be accessible only by computing devices that contribute to the computation (e.g., a first computing device 102 and at least one other computing device 106).

[0120] One or more other TEEs 104 may be associated with and/or implemented on other computing devices 106, according to the steps disclosed herein. Communication network 110 may include one or more wired and/or wireless networks. For example, communication network 110 may include a cellular network (e.g., a long-term evolution (LTE®) network, a third generation (3G) network, a fourth generation (4G) network, a code division multiple access (CDMA) network, and/or the like), a public land mobile network (PLMN), a local area network (LAN), a wide area network (WAN), a metropolitan area network (MAN), a telephone network (e.g., the public switched telephone network (PSTN)), a private network, an ad hoc network, a mesh network, a beacon network, an intranet, the Internet, a fiber optic-based network, a cloud computing network, and/or the like, and/or a combination of these or other types of networks.

[0121] In some non-limiting embodiments or aspects of multi-party computation between two parties, first computing device 102 may communicate a first computation input (e.g., the first party's contribution to the computation) and a first portion (e.g., share) of a one-time key to the TEE 104. TEE 104 may receive a second computation input (e.g., a second party's contribution to the computation) and a second portion (e.g., share) of the one-time key from at least one other computing device 106. TEE 104 may generate the one-time key from the first portion of the one-time key and the second portion of the one-time key. The TEE 104 may further generate the one-time key from a unique nonce associated with an instance of the computation (e.g., for each iteration of the computation when multiple successive computations are carried out between parties). The first portion of the one-time key may be generated from a private key (e.g., of a private-public key pair) associated with the first computing device 102. The second portion of the one-time key may be generated from a private key (e.g., of a private-public key pair) associated with the at least one other computing device 106. The first portion and the second portion of the one-time key may be communicated to the TEE 104 in a first time period and used by the TEE 104 to generate a plurality of one-time keys over successive instances of multi-party computation in subsequent time periods. TEE 104 may generate the computation based on the first computation input and the second computation input. TEE 104 may further encrypt the computation result using the one-time key to produce an encrypted output. First computing device 102 may receive the encrypted output from the TEE 104.

[0122] First computing device 102 may further communicate the encrypted output to at least one other computing device 106. In response to receiving the encrypted output, the at least one other computing device 106 may generate a digital signature indicating that the at least one other computing device 106 received the encrypted output. First computing device 102 may receive the digital signature from the at least one other computing device 106. In response to receiving the digital signature, first computing device 102 may communicate the first portion of the one-time key to the at least one other computing device 106. The first computing device 102 may then wait to receive the second portion of the one-time key from the at least one other computing device 106.

[0123] In response to receiving the second portion from the at least one other computing device 106, the first computing device 102 may determine the one-time key

from the first portion of the one-time key and the second portion of the one-time key. The first computing device 102 may further determine the computation by decrypting the encrypted output with the one-time key.

[0124] In response to not receiving the second portion of the one-time key from the at least one other computing device 106, the first computing device 102 may execute a fallback computation process using the TEE 104 and a ledger accessible to both the first computing device 102 and the at least one other computing device 106 to determine the computation.

[0125] In some non-limiting embodiments or aspects, the first computing device 102 may execute the fallback process in a series of steps. First computing device 102 may publish, to the shared ledger 108, the first portion of the one-time key and the digital signature received from the second computing device. The one-time key that is published to the shared ledger 108 may be encrypted with a public key associated with the at least one other computing device 106. First computing device 102 may receive a proof of publication from the ledger 108 and input the proof of publication to the TEE 104. TEE 104 may validate the proof of publication and, in response, send the one-time key, or the second portion of the one-time key, to the first computing device 102. First computing device 102 may receive the one-time key from the TEE 104, or the second portion of the one-time key with which the one-time key may be derived, and determine the computation by decrypting the encrypted output with the one-time key.

[0126] The number and arrangement of devices and networks shown in FIG. 1 are provided as an example. There may be additional devices and/or networks, fewer devices and/or networks, different devices and/or networks, or differently arranged devices and/or networks than those shown in FIG. 1. Furthermore, two or more devices shown in FIG. 1 may be implemented within a single device, or a single device shown in FIG. 1 may be implemented as multiple, distributed devices. Additionally or alternatively, a set of devices (e.g., one or more devices) of environment 100 may perform one or more functions described as being performed by another set of devices of environment 100.

[0127] Referring now to FIG. 2, illustrated is a diagram of example components of device 200. Device 200 may correspond to one or more devices of a first computing device 102, a trusted execution environment (TEE) 104, one or more other computing devices 106, a shared ledger 108, and/or a communication network 110. In some non-

limiting embodiments or aspects, one or more devices of the foregoing may include at least one device 200 and/or at least one component of device 200. As shown in FIG. 2, device 200 may include bus 202, processor 204, memory 206, storage component 208, input component 210, output component 212, and communication interface 214.

[0128] Bus 202 may include a component that permits communication among the components of device 200. In some non-limiting embodiments or aspects, processor 204 may be implemented in hardware, software, or a combination of hardware and software. For example, processor 204 may include a processor (e.g., a central processing unit (CPU), a graphics processing unit (GPU), an accelerated processing unit (APU), etc.), a microprocessor, a digital signal processor (DSP), and/or any processing component (e.g., a field-programmable gate array (FPGA), an application-specific integrated circuit (ASIC), etc.) that can be programmed to perform a function. Memory 206 may include random access memory (RAM), read-only memory (ROM), and/or another type of dynamic or static storage device (e.g., flash memory, magnetic memory, optical memory, etc.) that stores information and/or instructions for use by processor 204.

[0129] Storage component 208 may store information and/or software related to the operation and use of device 200. For example, storage component 208 may include a hard disk (e.g., a magnetic disk, an optical disk, a magneto-optic disk, a solid state disk, etc.), a compact disc (CD), a digital versatile disc (DVD), a floppy disk, a cartridge, a magnetic tape, and/or another type of computer-readable medium, along with a corresponding drive.

[0130] Input component 210 may include a component that permits device 200 to receive information, such as via user input (e.g., a touchscreen display, a keyboard, a keypad, a mouse, a button, a switch, a microphone, a camera, etc.). Additionally or alternatively, input component 210 may include a sensor for sensing information (e.g., a global positioning system (GPS) component, an accelerometer, a gyroscope, an actuator, etc.). Output component 212 may include a component that provides output information from device 200 (e.g., a display, a speaker, one or more light-emitting diodes (LEDs), etc.).

[0131] Communication interface 214 may include a transceiver-like component (e.g., a transceiver, a separate receiver and transmitter, etc.) that enables device 200 to communicate with other devices, such as via a wired connection, a wireless connection, or a combination of wired and wireless connections. Communication

interface 214 may permit device 200 to receive information from another device and/or provide information to another device. For example, communication interface 214 may include an Ethernet interface, an optical interface, a coaxial interface, an infrared interface, a radio frequency (RF) interface, a universal serial bus (USB) interface, a Wi-Fi® interface, a cellular network interface, and/or the like.

[0132] Device 200 may perform one or more processes described herein. Device 200 may perform these processes based on processor 204 executing software instructions stored by a computer-readable medium, such as memory 206 and/or storage component 208. A computer-readable medium (*e.g.*, a non-transitory computer-readable medium) is defined herein as a non-transitory memory device. A non-transitory memory device includes memory space located inside of a single physical storage device or memory space spread across multiple physical storage devices.

[0133] Software instructions may be read into memory 206 and/or storage component 208 from another computer-readable medium or from another device via communication interface 214. When executed, software instructions stored in memory 206 and/or storage component 208 may cause processor 204 to perform one or more processes described herein. Additionally or alternatively, hardwired circuitry may be used in place of or in combination with software instructions to perform one or more processes described herein. Thus, embodiments or aspects described herein are not limited to any specific combination of hardware circuitry and software.

[0134] Memory 206 and/or storage component 208 may include data storage or one or more data structures (*e.g.*, a database, and/or the like). Device 200 may be capable of receiving information from, storing information in, communicating information to, or searching information stored in the data storage or one or more data structures in memory 206 and/or storage component 208. For example, the information may include encryption data, input data, output data, transaction data, account data, or any combination thereof.

[0135] The number and arrangement of components shown in FIG. 2 are provided as an example. In some non-limiting embodiments or aspects, device 200 may include additional components, fewer components, different components, or differently arranged components than those shown in FIG. 2. Additionally or alternatively, a set of components (*e.g.*, one or more components) of device 200 may perform one or

more functions described as being performed by another set of components of device 200.

[0136] Referring now to FIG. 3, provided is a method 300 for secure real-time n-party computation, according to non-limiting embodiments or aspects. The steps of method 300 may be executed by a first computing device 102, a TEE 104, and/or another computing device 106. One or more steps of method 300 may be executed by a same or different computing device as another step of method 300. While the steps of method 300 are described with respect to a two-party computation, it will be appreciated that the principles and steps may be expanded to, and applicable to, higher numbers of parties as described herein.

[0137] In step 302, a first computation input and a first portion of a one-time key may be communicated. For example, a first computing device 102 may communicate, to a TEE 104 associated with the first computing device 102, a first computation input and a first portion of a one-time key. Step 302 may be triggered in response to input by a first party user in an interface of the first computing device 102. Step 302 may be procedurally and sequentially required before proceeding to step 304.

[0138] In step 304, an encrypted output may be received. For example, first computing device 102 may receive, from the TEE 104, an encrypted output of a computation based on the first computation input and a second computation input. The second computation input and a second portion of the one-time key may be communicated to the TEE 104 by a second computing device (e.g., one of the other computing devices 106) preceding step 304. The encrypted output may be encrypted by the TEE 104 with the one-time key generated from the first portion of the one-time key and the second portion of the one-time key. Step 304 may be procedurally and sequentially required before proceeding to step 306.

[0139] In step 306, the encrypted output may be communicated to the second computing device. For example, first computing device 102 may communicate the encrypted output (received from the TEE 104 in step 304) to the second computing device. Step 306 may be triggered in response to input by a first party user in an interface of the first computing device 102. Step 306 may be procedurally and sequentially required before proceeding to step 308.

[0140] In step 308, a digital signature may be received from the second computing device. For example, the first computing device 102 may receive a digital signature from the second computing device indicating that the second computing device

received the encrypted output in step 306. Step 308 may be procedurally and sequentially required before proceeding to step 310.

[0141] In step 310, the first portion of the one-time key may be communicated to the second computing device. For example, the first computing device 102 may communicate the first portion of the one-time key to the second computing device. If the second computing device sends its portion of the one-time key back in return, the first computing device 102 may proceed to determine the computation under a protocol where no fallback computation process is required (step 312). If the second computing device does not send its portion of the one-time key back in return, the first computing device 102 may proceed to determine the computation under a protocol where the fallback computation process is required (step 314). Step 310 may be procedurally and sequentially required before proceeding to step 312 or step 314.

[0142] In step 312, the one-time key may be determined and the encrypted output may be decrypted. For example, the first computing device 102 may, in response to receiving the second portion of the one-time key from the second computing device, determine the one-time key from the first portion of the one-time key and the second portion of the one-time key, and determine the computation by decrypting the encrypted output with the one-time key.

[0143] In step 314, the fallback computation process may be executed. For example, the first computing device 102 may, in response to not receiving the second portion of the one-time key from the second computing device, execute a fallback computation process using the TEE and a ledger accessible to both the first computing device 102 and the second computing device to determine the computation.

[0144] Referring now to FIG. 4, provided is a method 400 for secure real-time n-party computation, according to non-limiting embodiments or aspects. The steps of method 400 may be executed by a first computing device 102, a TEE 104, and/or another computing device 106. One or more steps of method 400 may be executed by a same or different computing device as another step of method 400. While the steps of method 400 are described with respect to a two-party computation, it will be appreciated that the principles and steps may be expanded to, and applicable to, higher numbers of parties as described herein. It will further be appreciated that method 400 provides an illustrative fallback computation process for use in step 314 of FIG. 3.

[0145] In step 402, the first portion of the one-time key and digital signature may be published to the ledger 108. For example, first computing device 102 may publish the first portion of the one-time key and digital signature received from the second computing device to the ledger 108. Step 402 may be procedurally and sequentially required before proceeding to step 404.

[0146] In step 404, a proof of publication may be received. For example, first computing device 102 may receive a proof of publication from the ledger 108 indicative of the first portion of the one-time key and the digital signature being published to the ledger 108. Step 404 may be procedurally and sequentially required before proceeding to step 406.

[0147] In step 406, the proof of publication may be input to the TEE 104. For example, first computing device 102 may input the proof of publication to the TEE 104. This will indicate to the TEE 104 that the first portion of the one-time key and the digital signature were published to the ledger 108. Step 406 may be procedurally and sequentially required before proceeding to step 408.

[0148] In step 408, the one-time key may be received from the TEE 104. For example, first computing device 102, in response to inputting the proof of publication to the TEE 104, may receive the one-time key from the TEE 104. It will also be appreciated that the first computing device 102 may receive the second portion of the one-time key, by which the one-time key may be determined in combination with the first portion of the one-time key. Step 408 may be procedurally and sequentially required before proceeding to step 410.

[0149] In step 410, the computation may be determined. For example, first computing device 102 may determine the computation by decrypting the encrypted output with the one-time key. In this manner, despite the second computing device aborting the process early and not sending the second portion of the one-time key to the first computing device 102, the first computing device 102 has a means of retrieving the one-time key as a whole in a manner that also makes sure the second computing device has access to the first portion of the one-time key, should the second computing device's departure from the process been unintentional, e.g., due to an error.

[0150] Referring to FIG. 5, depicted is a schematic diagram of a system 500 for secure, real-time, n-party computation, according to non-limiting embodiments or aspects. The system 500 of FIG. 5 depicts a multi-party computation of two parties: a

first party (P_1) 502a and a second party (P_2) 502b. The above-described two-party protocol may be employed in the system 500, as shown. Each party 502a, 502b is associated with its own party data 504a, 504b, which may be stored on and communicated by respective computing devices. For example, the first party 502a possesses a first set of party data 504a on a first computing device, and the second party 502b possesses a second set of party data 504b on a second computing device. The parties 502a, 502b work together to complete a joint computation 506 by communicating their respective party data 504a, 504b for use in the joint computation 506, and to receive the output of the joint computation 506 in an order of operations that protects honest parties should a dishonest party break the protocol early.

[0151] Referring to FIG. 6, depicted is a schematic diagram of a system 600 for secure, real-time, n-party computation, according to non-limiting embodiments or aspects. The system 600 of FIG. 6 depicts a multi-party computation of more than two parties: a first party (P_1) 502a, a second party (P_2) 502b, and any number of additional parties up to an nth party (P_N) 502n. The above-described n-party protocol may be employed in the system 600, as shown. Each party 502a, 502b, 502n is associated with its own party data 504a, 504b, 504n, which may be stored on and communicated by respective computing devices. For example, the first party 502a possesses a first set of party data 504a on a first computing device, the second party 502b possesses a second set of party data 504b on a second computing device, and any number of additional parties up to the nth party 502n may possess an nth set of party data 504n on an nth computing device. The parties 502a, 502b, 504n work together to complete a joint computation 606 by communicating their respective party data 504a, 504b, 504n for use in the joint computation 606, and to receive the output of the joint computation 606 in an order of operations that protects honest parties should a dishonest party break the protocol early.

[0152] Referring to FIG. 7, depicted is pseudocode for a method for secure real-time n-party computation, according to non-limiting embodiments or aspects. In particular, FIG. 7 depicts pseudocode for an enclave program for use in a two-party fair exchange protocol. The process depicted in FIG. 7 may be hosted and executed on a TEE 104.

[0153] Referring to FIGS. 8 and 9, depicted is pseudocode for a method for secure real-time n-party computation, according to non-limiting embodiments or aspects. In particular, FIGS. 8 and 9 depict pseudocode for a two-party fair exchange protocol.

The process depicted in FIGS. 8 and 9 may be executed by one or more computing devices, including a first computing device 102 associated with a first party and at least one other computing device 106 (e.g., a second computing device) associated with a second party.

[0154] Referring to FIG. 10, depicted is pseudocode for a method for secure real-time n-party computation, according to non-limiting embodiments or aspects. In particular, FIG. 10 depicts pseudocode for an enclave program for use in an n-party fair exchange protocol. The process depicted in FIG. 10 may be hosted and executed on one or more TEE 104.

[0155] Referring to FIGS. 11 and 12, depicted is pseudocode for a method for secure real-time n-party computation, according to non-limiting embodiments or aspects. In particular, FIGS. 11 and 12 depict pseudocode for an n-party fair exchange protocol. The process depicted in FIGS. 11 and 12 may be executed by one or more computing devices, including a first computing device 102 associated with a first party and a plurality of other computing devices 106 associated with other parties.

[0156] Referring to FIGS. 13 and 14, depicted is pseudocode for a method for secure real-time n-party computation, according to non-limiting embodiments or aspects. In particular, FIGS. 13 and 14 depict pseudocode for an n-party fair exchange protocol for a setting including multiple ledgers 108. The process depicted in FIGS. 13 and 14 may be executed by one or more computing devices, including a first computing device 102 associated with a first party and a plurality of other computing devices 106 associated with other parties.

[0157] Referring to FIG. 15, depicted is pseudocode for a method for secure real-time n-party computation, according to non-limiting embodiments or aspects. In particular, FIG. 15 depicts pseudocode for an enclave program for use in an n-party fair coin tossing protocol. The process depicted in FIG. 15 may be hosted and executed on one or more TEE 104.

[0158] Referring to FIGS. 16 and 17, depicted is pseudocode for a method for secure real-time n-party computation, according to non-limiting embodiments or aspects. In particular, FIGS. 16 and 17 depict pseudocode for an n-party fair coin tossing protocol. The process depicted in FIGS. 16 and 17 may be executed by one or more computing devices, including a first computing device 102 associated with a first party and a plurality of other computing devices 106 associated with other parties.

[0159] Although the disclosure has been described in detail for the purpose of illustration based on what is currently considered to be the most practical and non-limiting embodiments or aspects, it is to be understood that such detail is solely for that purpose and that the disclosure is not limited to the disclosed embodiments or aspects, but, on the contrary, is intended to cover modifications and equivalent arrangements that are within the spirit and scope of the appended claims. For example, it is to be understood that the present disclosure contemplates that, to the extent possible, one or more features of any embodiment or aspect can be combined with one or more features of any other embodiment or aspect.

WHAT IS CLAIMED IS

1. A system for multi-party computation between a first party and a second party, the system comprising:

a first computing device associated with the first party; and

a trusted execution environment (TEE) associated with the first computing device, wherein the first computing device is programmed or configured to:

communicate, to the TEE, a first computation input and a first portion of a one-time key;

receive, from the TEE, an encrypted output of a computation based on the first computation input and a second computation input, wherein the second computation input and a second portion of the one-time key are communicated to the TEE by a second computing device associated with the second party, and wherein the encrypted output is encrypted by the TEE with the one-time key;

communicate the encrypted output to the second computing device;

receive a digital signature from the second computing device indicating that the second computing device received the encrypted output;

communicate the first portion of the one-time key to the second computing device; and

(i) in response to receiving the second portion of the one-time key from the second computing device, determine the one-time key from the first portion of the one-time key and the second portion of the one-time key, and determine the computation by decrypting the encrypted output with the one-time key; or

(ii) in response to not receiving the second portion of the one-time key from the second computing device, execute a fallback computation process using the TEE and a ledger accessible to both the first computing device and the second computing device to determine the computation.

2. The system of claim 1, wherein the first computing device is further programmed or configured to, in the fallback computation process:

publish, to the ledger, the first portion of the one-time key and the digital signature received from the second computing device;

receive a proof of publication from the ledger;

input the proof of publication to the TEE;
receive the one-time key from the TEE; and
determine the computation by decrypting the encrypted output with the one-time key.

3. The system of claim 2, wherein the first portion of the one-time key that is published to the ledger by the first computing device is encrypted with a public key associated with the second computing device.

4. The system of claim 1, wherein the one-time key is generated by the TEE from the first portion of the one-time key, the second portion of the one-time key, and a unique nonce associated with an instance of the computation.

5. The system of claim 4, wherein the first portion of the one-time key is generated from a private key associated with the first computing device, and wherein the second portion of the one-time key is generated from a private key associated with the second computing device.

6. The system of claim 5, wherein the first portion of the one-time key and the second portion of the one-time key are communicated to the TEE in a first time period and are configured to generate a plurality of one-time keys over successive instances of multi-party computation in subsequent time periods.

7. The system of claim 1, wherein the ledger comprises an append-only blockchain and is configured to be accessible only by computing devices that contribute to the computation.

8. A computer program product for multi-party computation between a first party and a second party, the computer program product being stored on a first computing device associated with the first party and comprising at least one non-transitory computer-readable medium including program instructions that, when executed by at least one processor of the first computing device, cause the at least one processor to:

communicate, to a trusted execution environment (TEE) associated with the first computing device, a first computation input and a first portion of a one-time key;

receive, from the TEE, an encrypted output of a computation based on the first computation input and a second computation input, wherein the second computation input and a second portion of the one-time key are communicated to the TEE by a second computing device associated with the second party, and wherein the encrypted output is encrypted by the TEE with the one-time key;

communicate the encrypted output to the second computing device;

receive a digital signature from the second computing device indicating that the second computing device received the encrypted output;

communicate the first portion of the one-time key to the second computing device; and

(i) in response to receiving the second portion of the one-time key from the second computing device, determine the one-time key from the first portion of the one-time key and the second portion of the one-time key, and determine the computation by decrypting the encrypted output with the one-time key; or

(ii) in response to not receiving the second portion of the one-time key from the second computing device, execute a fallback computation process using the TEE and a ledger accessible to both the first computing device and the second computing device to determine the computation.

9. The computer program product of claim 8, wherein the program instructions further cause the at least one processor to, in the fallback computation process:

publish, to the ledger, the first portion of the one-time key and the digital signature received from the second computing device;

receive a proof of publication from the ledger;

input the proof of publication to the TEE;

receive the one-time key from the TEE; and

determine the computation by decrypting the encrypted output with the one-time key.

10. The computer program product of claim 9, wherein the first portion of the one-time key that is published to the ledger by the at least one processor is encrypted with a public key associated with the second computing device.

11. The computer program product of claim 8, wherein the one-time key is generated by the TEE from the first portion of the one-time key, the second portion of the one-time key, and a unique nonce associated with an instance of the computation.

12. The computer program product of claim 11, wherein the first portion of the one-time key is generated from a private key associated with the first computing device, and wherein the second portion of the one-time key is generated from a private key associated with the second computing device.

13. The computer program product of claim 12, wherein the first portion of the one-time key and the second portion of the one-time key are communicated to the TEE in a first time period and are configured to generate a plurality of one-time keys over successive instances of multi-party computation in subsequent time periods.

14. The computer program product of claim 8, wherein the ledger comprises an append-only blockchain and is configured to be accessible only by computing devices that contribute to the computation.

15. A computer-implemented method for multi-party computation between a first party and a second party, the method comprising:

Communicating, with at least one processor of a first computing device associated with the first party, to a trusted execution environment (TEE) associated with the first computing device, a first computation input and a first portion of a one-time key;

receiving, with at least one processor of the first computing device, from the TEE, an encrypted output of a computation based on the first computation input and a second computation input, wherein the second computation input and a second portion of the one-time key are communicated to the TEE by a second computing

device associated with the second party, and wherein the encrypted output is encrypted by the TEE with the one-time key;

communicating, with at least one processor of the first computing device, the encrypted output to the second computing device;

receiving, with at least one processor of the first computing device, a digital signature from the second computing device indicating that the second computing device received the encrypted output;

communicating, with at least one processor of the first computing device, the first portion of the one-time key to the second computing device; and

in response to not receiving the second portion of the one-time key from the second computing device, executing, with at least one processor of the first computing device, a fallback computation process using the TEE and a ledger accessible to both the first computing device and the second computing device to determine the computation.

16. The method of claim 15, wherein the fallback computation process comprises:

publishing, with at least one processor of the first computing device, to the ledger, the first portion of the one-time key and the digital signature received from the second computing device;

receiving, with at least one processor of the first computing device, a proof of publication from the ledger;

inputting, with at least one processor of the first computing device, the proof of publication to the TEE;

receiving, with at least one processor of the first computing device, the one-time key from the TEE; and

determining, with at least one processor of the first computing device, the computation by decrypting the encrypted output with the one-time key.

17. The method of claim 16, wherein the first portion of the one-time key that is published to the ledger is encrypted with a public key associated with the second computing device.

18. The method of claim 15, wherein the one-time key is generated by the TEE from the first portion of the one-time key, the second portion of the one-time key, and a unique nonce associated with an instance of the computation.

19. The method of claim 18, wherein the first portion of the one-time key is generated from a private key associated with the first computing device, and wherein the second portion of the one-time key is generated from a private key associated with the second computing device.

20. The method of claim 19, wherein the first portion of the one-time key and the second portion of the one-time key are communicated to the TEE in a first time period and are configured to generate a plurality of one-time keys over successive instances of multi-party computation in subsequent time periods.